computers are bad

2022-03-05 high definition radio

One of the great joys of the '00s was the tendency of marketers to apply the acronym "HD" to anything they possibly could. The funniest examples of this phenomenon are those where HD doesn't even stand for "High Definition," but instead for something a bit contrived like "Hybrid Digital." This is the case with HD Radio.

For those readers outside of these United States and Canada (actually Mexico as well), HD Radio might be a bit unfamiliar. In Europe, for example, a standard called DAB for Digital Audio Broadcasting is dominant and, relative to HD radio, highly successful. Another relatively widely used digital broadcast standard is Digital Radio Mondiale, confusingly abbreviated DRM, which is more widely used in the short and medium wave bands than in VHF where we find most commercial broadcasting today... but that's not a limitation, DRM can be used in the AM and FM broadcast bands.

HD radio differs from these standards in two important ways: first, it is intended to completely coexist with analog broadcasting due to the lack of North American appetite to eliminate analog. Second, no one uses it.

HD Radio broadcasts have been on the air in the US since the mid '00s. HD broadcasts are reasonably common now, with 9 HD radio carriers carrying 16 stations here in Albuquerque. Less common are HD radio receivers. Many, but not all, modern car stereos have HD Radio support. HD receivers outside of the car center console are vanishingly rare. Stereo receivers virtually never have HD decoding, and due to the small size of the market standalone receivers run surprisingly expensive. I am fairly comfortable calling HD Radio a failed technology in terms of its low adoption, but since it falls into the broader market of broadcast radio standards are low. We can expect HD Radio stations to remain available well into the future and continue to offer some odd programming.

Santa Fe's 104.1 KTEG ("The Edge"), for example, a run of the mill iHeartMedia alt rock station, features as its HD2 "subcarrier" a station called Dance Nation '90s. The clearly automated programming includes Haddaway's "What Is Love" seemingly every 30 minutes and no advertising whatsoever, because it clearly doesn't have enough listeners for any advertisers to be willing to pay for it. And yet it keeps on broadcasting, presumably an effort by iHeartMedia to meet programming diversity requirements while still holding multiple top-40 licenses in the Albuquerque-Santa Fe market region [1].

So what is all this HD radio stuff? What is a subcarrier? And just what makes HD radio "Hybrid Digital?" HD Radio has gotten some press lately because of the curious failure mode of some Mazda head units, and that's more attention than it's gotten for years, so let's look a bit at the details.

First, HD Radio is primarily, in the US, used in a format called In-Band On-Channel, or IBOC. The basic idea is that a conventional analog radio station continues to broadcast while an HD Radio station is superimposed on the same frequency. The HD Radio signal is found "outside" of the analog signal, as two prominent sideband signals outside of the bandwidth of analog FM stereo.

While the IBOC arrangement strongly resembles a single signal with both analog and digital components, in practice it's very common for the HD signal to be broadcast by a separate transmitter and antenna placed near the analog transmitter (in order to minimize destructive interference issues). This isn't quite considered the "correct" implementation but is often cheaper since it avoids the need to make significant changes to the existing FM broadcast equipment... which is often surprisingly old.

It's completely possible for a radio station to transmit *only* an HD signal, but because of the rarity of HD receivers this has not become popular. The FCC does not normally permit it, and has declined to extend the few experimental licenses that were issued for digital-only operation. As a result, we see HD Radio basically purely in the form of IBOC. More commonly, HD Radio supports both a full hybrid mode with conventional FM audio clarity and also an "extended" mode in which the digital sidebands intrude on the conventional FM bandwidth. This results in mono-only, reduced-quality FM audio, but allows for a greater digital data rate.

HD Radio was developed and continues to be maintained by a company called iBiquity, which was acquired by DTS, which was acquired by Xperia. iBiquity maintains a patent pool and performs (minimal) continuing development on the standard. iBiquity makes their revenue from a substantial up-front license fee for radio stations to use HD Radio, and from royalties on revenue from subcarriers. To encourage adoption, no royalties are charged on each radio station's primary audio feed. Further encouraging adoption (although not particularly successfully), no royalty or license fees are required to manufacture HD Radio receivers.

The adoption of HD Radio in North America stems from an evaluation process conducted by the FCC in which several commercial options were considered. The other major competitor was FMeXtra, a generally similar design that was not selected by the FCC and so languished. Because US band planning for broadcast radio is significantly different from the European approach, DAB was not a serious contender (it has significant limitations due to the very narrow RF bandwidth available in Europe, a non-issue in the US where each FM radio station was effectively allocated 200kHz).

Layer 1

The actual HD Radio protocol is known more properly as NRSC-5, for its standards number issued by the National Radio Systems Council. The actual NRSC-5 protocol differs somewhat depending on whether the station is AM or FM (the widely different bandwidth characteristics of the two bands require different digital encoding approaches). In the more common case of FM, NRSC-5 consists of a set of separate OFDM data carriers, each conveying part of several logical channels which we will discuss later. A total of 18 OFDM subcarriers are typically present, plus several "reference" subcarriers which are used by receivers to detect and cancel certain types of interference.

If you are not familiar with OFDM or Orthogonal Frequency Division Multiplexing, it is an increasingly common encoding technique that essentially uses multiple parallel

digital signals (as we see with the 18 subcarriers in the case of HD Radio) to allow each individual signal to operate at a lower symbol rate. This has a number of advantages, but perhaps the most important is that it is typically used to enable the addition of a "guard interval" between each symbol. This intentional quiet period avoids subsequent symbols "blurring together" in the form of inter-symbol interference, a common problem with broadcast radio systems where multipath effects result in the same signal arriving multiple times at slight time offsets.

A variety of methods are used to encode the logical channels onto the OFDM subcarriers, things like scrambling and convolutional coding that improve the ability of receivers to recover the signal due to mathematics that I am far from an expert on. The end result is that an NRSC-5 standard IBOC signal in the FM band can convey somewhere from 50kbps to 150kbps depending on the operator's desired tradeoffs of bitrate to power and range.

The logical channels are the interface from layer 1 of NRSC-5 to layer 2. The number and type of logical channels depends on the band (FM or AM), the waveform (hybrid analog and digital, analog with reduced bandwidth and digital, or digital only), and finally the service mode, which is basically a configuration option that allows operators to select how digital capacity is allocated.

In the case of FM, five logical channels are supported... but not all at once. A typical full hybrid station broadcasts only primary channel P1 and a the PIDS channel, a low-bitrate channel for station identification. P1 operates at approximately 98kbps. For stations using an "extended" waveform with mono FM, the operator can select from configurations that provide 2-3 logical channels with a total bitrate of 110kbps to 148kbps. Finally, all-digital stations can operate in any extended service mode or at lower bitrates with different primary channels present. Perhaps most importantly, all-digital stations can include various combinations of *secondary* logical channels which can carry yet more data.

The curious system of primary channels is one that was designed basically to ease hardware implementation and is not very intuitive... we must remember that when NRSC-5 was designed, embedded computing was significantly more limited. Demodulation and decoding would have to be implemented in ASICs, and so many aspects of the protocol were designed to ease that process. At this point it is only important to understand that HD Radio's layer 1 can carry some combination of 4 primary channels along with the PIDS channel, which is very low bitrate but considered part of the primary channel feature set.

Layer 1, in summary, takes some combination of primary channels, the low-bitrate PIDS channel, and possibly several secondary channels (only in the case of all-digital stations) and encodes them across a set of OFDM subcarriers arranged just outside of the FM audio bandwidth. The design of the OFDM encoding and other features of layer 1 aid receivers in detecting and decoding this data.

Layer 2

Layer 2 operates on protocol data units or PDUs, effectively the packet of NRSC-5. Specifically, it receives PDUs from services and then distributes them to the layer 1 logical channels.

The services supported by NRSC are the Main Program Service or MPS which can carry

both audio (MPSA) and data (MPSD), the similar Supplemental Program Service which also conveys audio and data, the Advanced Application Service (AAS), and the Station Information Service (SIS).

MPS and SPS are where most of HD Radio happens. Each carries a program audio stream along with program data that is related to the audio stream——things like metadata of the currently playing track. These streams can go onto any logical channel at layer 1, depending on the bitrate required and available. An MPS stream is mandatory for an HD radio station, while an SPS is optional.

AAS is an optional feature that can be used for a variety of different purposes, mostly various types of datacasting, which we'll examine later. And finally, the SIS is the simplest of these services, as it has a dedicated channel at layer 1 (the PIDS previously mentioned). As a result, layer 2 just takes SIS PDUs and puts them directly on the layer 1 channel dedicated to them.

The most interesting part of layer 2 is the way that it muxes content together. Rather than sending PDUs for each stream, NRSC-5 will combine multiple streams within PDUs. This means that a PDU may contain only MPS or SPS audio, or it might contain some combination of MPS or SPS with other types of data. While this seems complicated, it has some convenient simplifying properties: PDUs can be emitted for each program stream at a fixed rate based on the audio codec rate. Any unused space in each PDU can then be used to send other types of data, such as for AAS, on an as-available basis. The situation is somewhat simplified for the receiver since it knows exactly when to expect PDUs containing program audio, and that program audio is always the start of a PDU.

The MPS and one or more SPS streams, if present, are not combined together but instead remain separate PDUs and are allocated to the logical channels in one of several fixed schemes depending on the number of SPS present and the broadcast configuration used by the station. In the most common configuration, that of one logical channel on a full hybrid radio station, the MPS and up to two SPS are multiplexed onto the single logical channel. In more complex scenarios such as all-digital stations, the MPS and three SPS may be multiplexed across three logical channels. Conceptually, up to seven distinct SPS identified by a header field can be supported, although I'm not aware of anyone actually implementing this.

It is worth discussing here some of the practical considerations around the MPS and SPS. NRSC-5 requires that an MPS always be present, and the MPS must convey a "program 0" which cannot be stopped and started. This is the main audio channel on an HD radio station. The SPS, though, are used to convey "subcarrier" [2] stations. This is the capability behind the "HD2" second audio channel present on some HD radio stations, and it's possible, although not at all common, to have an HD3 or even HD4.

Interestingly, the PDU "header" is not placed at the beginning of the PDU. Instead, its 24-bit sequence (chosen off a list based on what data types are present in the PDU) are interleaved throughout the body of the PDU. This is intended to improve robustness by allowing the receiver to correctly determine the PDU type even when only part of the PDU is received. PDUs always contain mixed data in a fixed order (program data, opportunistic data, fixed data), with a "data delimiter" sequence after the program audio and a fixed data length value placed at the end. This assists receivers in interpreting any partial PDUs, since they can "backtrack" from the length suffix to identify the full fixed data section and then search further back for the "data delimiter" to identify the full opportunistic data section.

And that's layer 2: audio, opportunistic data, and fixed data are collected for the MPS and any SPS and/or AAS, gathered into PDUs, and then sent to layer 1 for transmission. SID is forwarded directly to layer 1 unmodified.

Applications

NRSC–5's application layer runs on top of layer 2. Applications consist most obviously of the MPS and SPS streams, which are used mainly to convey audio... you know, the thing that a radio station does. This can be called the Audio Transport application and it runs the same way whether producing MPS (remember, this is the main audio feed) or SPS (secondary audio feeds or subcarriers).

Audio transport starts with an audio encoder, which is a proprietary design called HDC or High–Definition Coding. HDC is a DCT–based lossy compression algorithm which is similar to AAC but adjusted to have some useful properties for radio. Among them, HDC receives audio data (as PCM) at a fixed rate and then emits encoded blocks at a fixed rate––but variable size. This variable size but fixed rate is convenient to receivers but also makes "opportunistic data," as discussed earlier, possible, because many PDUs will have spare room at the end.

Another useful feature of HDC is its multi–stream output. HDC can be configured to produce two different bit streams, a "core" bit stream which is lower bitrate but sufficient to reproduce the audio at reduced quality, and an "enhanced" data stream that allows the reproduction of higher fidelity audio. The core bit stream can be placed on a different layer 1 channel than the enhanced data stream, allowing receivers to decode only one channel and still produce useful audio when the second channel is not available due to poor reception quality. This is not typically used by hybrid stations, instead it's a feature intended for extended and digital–only stations.

The variable size of the audio data and variable size of PDUs creates some complexity for receivers, so the audio transport includes some extra data about sample rate and size to assist receivers in selecting an appropriate amount of buffering to ensure that the program audio does not underrun despite bursts of large audio samples and fixed data. This results in a fixed latency from encoding to decoding, which is fairly short but still a bit behind analog radio. This latency is sometimes apparent on receivers that attempt to automatically select between analog and digital signals, even though stations *should* delay their analog audio to match the NRSC–5 encoder.

Finally, the audio transport section of each PDU (that is, the MPS or SPS part at the beginning) contains regular CRC checksums that are used by the receiver to ensure that any bad audio data is discarded rather than decoded.

MPS and SPS audio is supplemented by Program Service Data (PSD), which can be either associated with the MPS (MPSD) or an SPS (SPSD). The PSD protocol generates PDUs which are provided to the audio transport to be incorporated into audio PDUs at the very beginning of the MPS or SPS data. The PSD is rather low bitrate as it receives only a small number of bytes in each PDU. This is quite sufficient, as the PSD only serves to move small, textual metadata about the audio. Most commonly this is the title, artist, and album, although a few other fields are included as well such as structured metadata for advertisements, including a field for price of the advertised deal. This feature is rarely, if ever, used.

The PSD data is transmitted continuously in a loop, so that a receiver that has just tuned to a station can quickly decode the PSD and display information about whatever is being broadcast. The looping PSD data changes whenever required, typically based on an outside system (such as a radio automation system) sending new PSD data to the NRSC-5 encoder over a network connection. PSD data is limited to 1024 bytes total and, as a minimum, the NRSC-5 specification requires that the title and artist fields be populated. Oddly, it makes a half-exception for cases where no information on the audio program is available: the artist field can be left empty, but the title field must be populated with some fixed string. Some radio stations have added an NRSC-5 broadcast but not upgraded their radio automation to provide PSD data to the encoder; in this case it's common to transmit the station call sign or name as the track title, much as is the case with FM Radio Data Service.

Interestingly, the PSD data is viewed as a set of ID3 tags and, even though very few ID3 fields are supported, it is expected that those fields be in the correct ID3 format including version prefixes.

Perhaps the most sophisticated feature of NRSC-5 is the Advanced Application Service transport or AAS. AAS is a flexible system intended to send just about any data alongside the audio programs. Along with PDUs, the audio transport generates a metadata stream indicating the length of the PDUs which is used. The AAS can use that value to determine how many bytes are free, and then fill them with opportunistic data of whatever type it likes. As a result, the AAS basically takes advantage of any "slack" in the radio broadcast's capacity, as well as reserving a portion for fixed data if desired by the station operator.

AAS data is encoded into AAS packets, an organizational unit independent of PDUs (and included within PDUs generated by the audio transport) and loosely based on computer networking conventions. Interestingly, AAS packets may be fragmented or combined to fit into available space in PDUs. To account for this variable structure, AAS specifies a transport layer below AAS packets which is based on HDLC (ISO high-level data link control) or PPP (point-to-point protocol, which is closely related to HDLC and very similar). So, in a way, AAS consists of a loosely computer-network-like protocol over a protocol roughly based on PPP over audio transport PDUs over OFDM.

Each AAS packet header specifies a sequence number for reconstruction of large payloads and a port number, which indicates to the receiver how it should handle the packet (or perhaps instead ignore the packet). A few ranges of port numbers are defined, but the vast majority are left to user applications.

Port numbers are two bytes, and so there's a large number of applications possible. Very few are defined by specification, limited basically to port numbers for supplemental PSD. This might be a bit confusing since PSD has its own reserved spot at the beginning of the audio transport. The PSD protocol itself is limited to only small amounts of text, and so when desired AAS can be used to send larger PSD-type payloads. The most common application of this "extra PSD" is album art, which can be sent as a JPG or PNG file in the AAS stream. In fact, multiple ports are reserved for each of MPSD (main PDS) and SPSD, allowing different types of extra data to be sent via AAS.

Ultimately, the AAS specification is rather thin... because AAS is a highly flexible feature that can be used in a number of ways. For example, AAS forms the basis of the Artist Experience service which allows for delivery of more complete metadata on musical tracks including album art. AAS can be used as the basis of almost any datacasting application, and is applied to everything from live traffic data to

distribution of educational material to rural areas.

Finally, in our tour of applications, we should consider the station information service or SIS. SIS is a very basic feature of NRSC-5 that allows a station to broadcast its identification (call sign and name) along with some basic services like a textual message related to the station and emergency alert system messages. SIS has come up somewhat repeatedly here because it receives special treatment; SIS is a very simple transport at a low bitrate and has its own dedicated logical channel for easy decoding. As a result, SIS PDUs are typically the first thing a receiver attempts to decode, and are very short and simple in structure.

To sum up the structure of HD radio, it is perhaps useful to look at it as a flow process: SID data is generated by the encoder and sent to layer 2 which passes it directly to layer 1, where it is transmitted on its own logical channel. PSD data is provided to the audio transport which embeds it at the beginning of audio PDUs. The audio transport informs the AAS encoder of the amount of available free space in a PDU, and the AAS encoder provides an appropriate amount of data to the audio transport to be added at the end of the PDU. This PDU is then passed to layer 2 which encapsulates it in a complete NRSC-5 PDU and arranges it into logical channels which are passed to layer 1. Layer 1 encodes the data into multiple OFDM carriers using a somewhat complex scheme that produces a digital signal that is easy for receivers to recover.

Non-Audio Applications of NRSC-5

While the NRSC-5 specification is clearly built mostly around transporting the main and secondary program audio, the flexibility of its data components like PSD and AAS allows its use for purposes other than audio. As a very simple example, SIS packets include a value called the "absolute local frame number" or ALFN that is effectively a timestamp, useful for receivers to establish the currency of emergency alert messages and for various data applications. Because the current time can be easily calculated from the ALFN, it can be used to set the clocks on HD radio receivers such as car head units. To support this, standard SIS fields include information on local time zone, daylight savings time, and even upcoming leap seconds.

SIS packets include a one-bit flag that indicates whether or not the ALFN is being generated based on a GPS-locked time source, or based on the NRSC-5 encoder's internal clock only. To avoid automatically adjusting radio clocks to an incorrect time (something that had plagued the earlier CEA protocol for automatic setting of VCR clocks via PBS member stations), NRSC-5 dictates that receivers must not set their display time based on a radio station's ALFN unless the flag indicating GPS lock is set. Unfortunately, it seems that it's rather uncommon for radio stations to equip their encoder with a GPS time source, and so in the Albuquerque market at least HD Radio-based automatic time setting does not work.

Other supplemental applications were included in the basic SIS as well, notably emergency alert messages. HD Radio stations can transmit emergency alert messages in text format with start and end times. In practice this seems to be appreciably less successful than the more flexible capability of SiriusXM, and ironically despite its cost to the consumer SiriusXM might have better market penetration than HD Radio.

NRSC-5's data capabilities can be used to deliver an enhanced metadata experience around the audio programming. The most significant implementation of this concept is

the "artist experience" service, a non-NRSC-5 standard promulgated by the HD Radio alliance that uses the AAS to distribute more extensive metadata including album art in image format. This is an appreciably more complex process and so is basically expected to be implemented in software on a general-purpose embedded operating system, rather than the hardware-driven decoding of audio programming and basic metadata. Of course this greater complexity lead more or less directly to the recent incident with Mazda HD radio receivers in Seattle, triggered by a station inadvertently transmitting invalid Artist Experience data in a way that seems to have caused the Mazda infotainment system to crash during parsing. Fortunately infotainment-type HD radio receivers typically store HD Radio metadata in nonvolatile memory to improve startup time when tuning to a station, so these Mazda receivers apparently repeatedly crashed every time they were powered on to such a degree that it was not possible to change stations (and avoid parsing the cached invalid file). Neat.

Since Artist Experience just sends JPG or PNG files of album art, we know that AAS can be used to transmit files in general (and looking at the AAS protocol you can probably easily come up with a scheme to do so). This opens the door to "datacasting," or the use of broadcast technology to distribute computer data. I have written on this topic before.

To cover the elements specific to our topic, New Mexico's KANW and some other public radio stations are experimenting with transmitting educational materials from local school districts as part of the AAS data stream on their HD2 subcarrier. Inexpensive dedicated receivers collect these files over time and store them on an SD card. These receiver devices also act as WiFi APs and offer the stored contents via an embedded web server. This allows the substantial population of individuals with phones, tablets, or laptops but no home internet or cellular service to retrieve their distance education materials at home, without having to drive into town for cellular service (the existing practice in many parts of the Navajo Nation, for example) [3].

There is potential to use HD Radio to broadcast traffic information services, weather information, and other types of data useful to car navigation systems. While there's a long history of datacasting this kind of information via radio, it was never especially successful and the need has mostly been obsoleted by ubiquitous LTE connectivity. In any case, the enduring market for this type of service (over-the-road truckers for example) has a very high level of SiriusXM penetration and so already receives this type of data.

Fall of the House of Hybrid Digital

In fact, the satellite angle is too big to ignore in an overall discussion of HD Radio. Satellite radio was introduced to the US at much the same time as HD Radio, although XM proper was on the market slightly earlier. Satellite has the significant downside of a monthly subscription fee. However, time seems to have shown that the meaningful market for enhanced broadcast radio consists mostly of people who are perfectly willing to pay a $20/mo subscription for a meaningful better service. Moreover, it consists heavily of people involved in the transportation industry (Americans listen to the radio basically only in vehicles, so it makes sense that the most dedicated radio listeners are those who spend many hours in motion). Since many of these people regularly travel across state lines, a nationwide service is considerably more useful than one where they have to hunt for a new good station to listen to as they pass through each urban area.

All in all, HD radio is not really competitive for today's serious radio listeners because it fails to address their biggest complaint, that radio is too local. Moreover, SiriusXM's ongoing subscription revenue seems to provide a much stronger incentive to quality than iHeartMedia's declining advertising relationships [4]. The result is that, for the most part, the quality of SiriusXM programming is noticeably better than most commercial radio stations, giving it a further edge over HD Radio.

Perhaps HD Radio is simply a case of poor product-market fit, SiriusXM having solved essentially the same problems but much better. Perhaps the decline of broadcast media never really gave it a chance. The technology is quite interesting, but adoption is essentially limited to car stereos, and not even that many of them. I suppose that's the problem with broadcast radio in general, though.

[1] The details here are complex and deserve their own post, but as a general idea the FCC attempts to maintain a diversity of radio programming in each market by refusing licenses to stations proposing a format that is already used by other stations. Unfortunately there are relatively few radio formats that are profitable to operate, so the broadcasting conglomerates tend to end up playing games with operating stations in minor formats, at little profit, in order to argue to the FCC that enough programming diversity is available to justify another top 40 or "urban" station.

[2] The term "subcarrier" is used this way basically for historical reasons and doesn't really make any technical sense. It's better to think of "HD2" as being a subchannel or secondary channel, but because of the long history of radio stations using *actual* subcarrier methods to convey an alternate audio stream the subcarrier term is stuck.

[3] It seems inevitable that, as has frequently happened in the history of datacasting, improving internet access technology will eventually obsolete this concept. I would strongly caution you against thinking this has already happened, though: even ignoring the issue of the long and somewhat undefined wait, Starlink is considerably more expensive than the typical rates for rural internet service in New Mexico. It is to some extent a false dichotomy to say that Starlink is cost uncompetitive with DSL considering that it can service a much greater area. However, I think a lot of "city folk" are used to the over-$100-per-month pricing typical of urban gigabit service and so view Starlink as inexpensive. They do not realize that, for all the downsides of rural DSL, it is very cheap. This reflects the tight budget of its consumers. For those who have access, CenturyLink DSL in New Mexico ranch country is typically $45/mo no-contract with no install fee and many customers use a $10/mo subsidized rate for low income households. Starlink's $99/mo and $500 initial is simply unaffordable in this market, especially since those outside of the CenturyLink service area have, on average, an even lower disposable income than those clustered near towns and highways.

[4] It is hard for me not to feel like iHeartMedia brought this upon themselves. They gained essentially complete control of the radio industry (with only even sadder Cumulus as a major competitor) and then squeezed it for revenue until US commercial radio programming had become, essentially, a joke. Modern commercial radio stations run on exceptionally tight budgets that have mostly eliminated any type of advantage they might have had due to their locality. This is most painfully apparent when you hear an iHeartMedia station give a rare traffic update (they seem to view this today as a mostly pro forma activity and do it as little as possible) in which the announcer pronounces "Montaño" and, more puzzlingly, "Coors" wrong in the span of a single sentence. I have heard a rumor that all of the iHeartMedia traffic announcements are done centrally from perhaps Salt Lake City but I do not know if this is true.